BIOLOGICALLY INSPIRED COMPUTATION

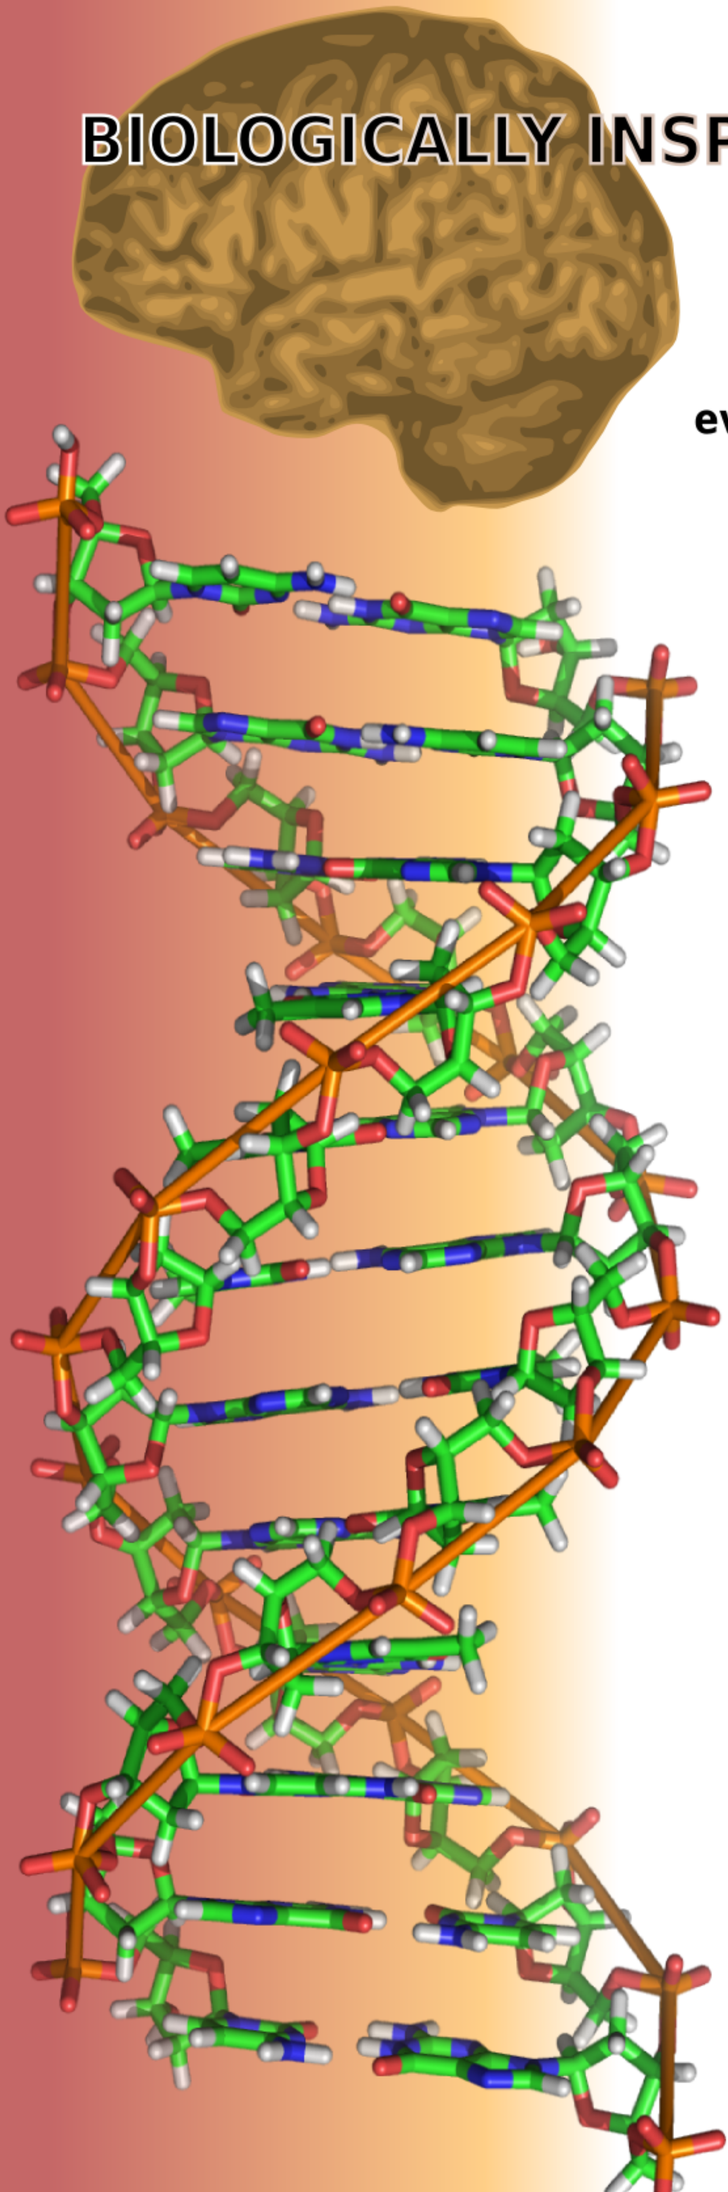

Emergence of Intelligence through evolving Artificial Neural Networks

Bachelor's Thesis

Erik de Bruijn
ANR 23.99.45
E.deBruijn@TilburgUniversity.nl

Supervised by:
Dr. H. Weigand
Dept. of Information Management

August 2007

VISIT…

LANZAROTE Caliente.COM

Abstract

Nature has evolved an abundance of valuable and inspirational examples. It has even given birth to intelligent species. The concept of intelligence is valuable to us and deserves further exploration. The 'brain processes' underlying intelligence are still little understood. Luckily, complex and often complicated phenomena have underlying principles that are not very complicated by themselves. In this thesis, (1) the principles for this evolutionary process were clarified and (2) evaluated for artificial application. (3) Finally we address whether intelligence might emerge. This thesis provides an overview of AI from biology to artificial implementation.

Ad 1) General principles thought to be essential are evolution, evolvable structures (substrates) and interaction with a rich and challenging environment. Specifically, neuronal structures have been essential to natural evolution of intelligence.

Ad 2) Both neuronal structures and evolution have been implemented artificially and have been combined, referred to as Artificial Neural Networks (ANNs) and Evolutionary Algorithms (EAs). Two experimental implementations are discussed and related to the theory. Evolution of virtual creatures' shapes to ANNs and EA and an artificial developing humanoid 'baby robot' to developmental psychology. Implementation challenges and issues are discussed, scaling and interconnection problems. Possible solutions are use of FPGA, aVLSI, neuromorphic engineering, optic-holographic and molecular computing devices.

Ad 3) Do the implementations have what is needed for intelligence to emerge? Will intelligence eventually arise? A theoretical computer science perspective and the view of Alastair Channon are presented. Gödel's incompleteness theorem and Searle's 'Chinese room' experiment are introduced. Channon argues that since we're unable to specify precisely what intelligence is, we should not expect it to emerge when using the fitness function in the traditional sense. Instead, Channon and others propose a co-evolution based approach.

Contents

Acknowledgments

I am grateful to my supervisor, Dr. Hans Weigand, for his support, patience and also for pressing me to simplify. As it is now, I hope there is much value in both the breath and depth of this work.

Without the loving support and faith of Esther, my parents and brother, this work would never have been accomplished. I'd like to thank my best of friends, Marius, Aike and Serge for their support and many inspirational discussions and accepting my dedication to this research project for over half a year. Furthermore, I am grateful to "cwenner" and "obsethryl" on `#ai` and `#neuroscience` @ irc.freenode.net for reviewing my work and suggesting articles. Also, I'd like to thank nature for making me inquisitive and eager to learn.

Chapter 1

Introduction

1.1 Background

Understanding the workings of nature is a fundamental pursuit of most sciences. Thanks to its vast diversity, it has offered us many examples and baffled us with its ingenuity.

It is striking to see that mechanisms that are in essence quite simple, can give birth to intelligent beings like ourselves. While the basic behavior and function of a single neuron or the process of evolution can be understood with relative ease, understanding the intelligence that it can produce has proved to be far more difficult.

The fact that we have the ability to consciously think about the world around us, suggests that we have enormous (symbolic) processing capacities. Human understanding goes far beyond the learning behavior of most (if not all) living organisms. The ability to learn has served as a stepping stone towards comprehension. This ability itself is one that has developed over many generations, species and millennia. The general tendency is that higher species (which exhibit more complex behavior) are more recent products of evolutions [MSS95].

The pursuit to understand nature has produced quantum physics, which in turn made possible innovations such as microprocessors. These microprocessors have also proved their enormous, ever growing, processing capacity. While modern microprocessors are also increasingly complex, they too are based on a relatively simple concept: the transistor. In essence, a transistor is quite similar to a neuron. The way in which these two building blocks are organized, however, is quite different.

The dominant system architecture for microprocessors is 'von Neumann', the common substrate silicon and the routing architecture VLSI[1]. The computational core of a microprocessor goes through a serial process of fetching and executing instructions at speeds of multiple GHz. A decade ago clock

[1]**V**ery **L**arge **S**cale **I**ntegration

frequency was a factor of a 100 slower (measured in MHz). Also, the amount of transistors on a chip has increased exponentially, according than Moore's famous law [Moo65]. The quantitative aspects of the physical structure (or physiology) of microchips is immense. Ray Kurzweil, AI researcher, inventor and analyst of technology trends, predicts that "[s]upercomputers will achieve one human brain capacity by 2010, and personal computers will do so by about 2020" [Kur99, Kur90].

Most models of brain organization consider nerve cells and their connections to be the brain's fundamental units of information processing. However, profoundly complex and intelligent activities occur within nerve cells [Ham03, pp. 9-10].

The first neural models developed by the psychologist Rosenblatt are highly similar to those developed in electrical engineering [Fau94, p. 23]. Now, ANNs and evolution are approached from the field of computing, besides biology, psychology, complexity and information theory. This highlights the interdisciplinary nature of these biologically inspired fields.

The architecture of the nervous system in humans is an enormous quantity of neurons working in parallel. Frequency of 'calculations' per neuron are very modest, at most 200 Hz [KW01, p. 129][2], but the brain makes up for this because of the parallelism: up to 10^{11} neurons are actively processing information at any moment [PWB87, p. 4]. Action potentials speeds range from 0.55 m/s and up to 120 m/s[3].

A topic on the application of biologically inspired principles to computing is necessarily multidisciplinary of nature. The questions that are addressed in this thesis could not be answered without drawing from a huge body of research spanning evolutionary biology, complexity theory, neurology, neurophysiology, evolutionary psychology[4], theoretical and applied computer science, Artificial Intelligence (AI), robotics and electronic engineering. Creation of AI requires the insights of many of these fields.

1.2 Research questions

The following main research questions will be addressed:

1. **What are driving principles for intelligence in biological systems, and what are their roles?** (chapter 2)

2. **Can we implement these principles artificially?** (chapter 3)

3. **Do they give rise to intelligence?** (chapter 4)

[2] McCulloch in 1952: "perhaps 100 bits/s", Kurzweil: 100Hz.

[3] 100 m/s [Ham03, p. 61], 111 m/s [EA03, p. 1], perhaps up to 120 m/s [KW01, p. 132], 120 m/s [Pen90]

[4] The study of behavior that uses principles of natural selection to account for human behaviors.[KW01]

Chapter 2 will deal with the first question, by exploring biological views on evolution and the brain. Principles such as evolution and neuronal structures will be identified as essential to intelligent.

The second question, 'Can we implement these driving principles artificially?', calls for answering a couple of subsequent questions:

1. What biologically-inspired systems have already been built?
2. To which degree are they biologically realistic?
3. What are currently barriers to implementation?

The third major question of this research (chapter 4), necessarily a philosophically loaded question, will not be answered conclusively. Instead, two views on the possibilities and limitations of evolution and the mind are discussed.

This thesis is a review of the literature concerning the above questions. Literature was selected based on frequencies of citation and general popularity and availability of criticism. An exception is the experiment on evolving morphology of virtual creatures by Goldstein, it is not well known. I have repeated the experiment and the experiment is similar the more widely known work of Karl Sims[5]. In the relatively small experiment, Goldstein brings forward an important lesson learned that I wish to present.

1.3 Motivation and rationale

1.3.1 Intelligence defined

" Intelligence usually means
'the ability to solve hard problems'. "
`Marvin Minsky`[6]

"The true sign of intelligence is not knowledge
but imagination."
`Albert Einstein`

"Intelligence is the ability to adapt to change."
`Stephen Hawking`

[5]Extensive material on Sims' virtual creatures can be found here: http://www.genarts.com/karl/

[6]Marvin Minsky has made many contributions to AI, cognitive psychology, mathematics, computational linguistics, robotics, and optics. In recent years he has worked chiefly on imparting to machines the human capacity for common sense reasoning. His conception of human intellectual structure and function is presented in his book "The Society of Mind" and further elaborated upon in his upcoming book "The Emotion Machine." [from www.KurzweilAI.net]. Quote seems to be from 1985.

"Viewed narrowly, there seem to be
almost as many definitions of intelligence
as there were experts asked to define it."
`R. J. Sternberg` quoted in [Gre98].

There exists no single, clear and concise definition of intelligence, however many have been suggested (see above and [LH07, Got97] for 70 more). The fact that the definition of intelligence has received such devoted attention and still resists clarification so thoroughly, is alarming[7] [Reb95, Jen99, part III, §11]. Because the first definitions given were rooted in psychology, specifically psychometrics (dealing mostly with measuring the human IQ), it inherited a cultural subjectivity of what is to be dubbed intelligent. Consequently, measuring how intelligent something is may be debated endlessly. Since the purpose of concise definitions is to prevent ambiguities and the resulting endless discussions (see your dictionary for a definition of the word 'definition'), we may want to lessen the definitional burden.

What intelligence involves is quite simple to identify: abstraction, learning and dealing with novelty. Whether it is exactly this set of abilities that is fundamental, or another is hard to determine. Certainly, there is some consensus around the data on intelligence. It shows that most suggested abilities are significantly correlated, leading to the discovery of the g factor[8]. Also, intelligence is graded, there is a smooth transition between systems, which everyone would agree to be not intelligent and truly intelligent systems [Hut04, p. 2]. If it is ever pin-pointed what intelligence precisely is, or when scientists finally concur, most theories that have strong support will all be very close because of this strong correlation and overlap of these notions. It is exactly the correlation that reveals that a general principle is responsible for emergence of these abilities. Moreover, Jensen emphasizes that "[i]t is not essentially a psychological or behavioral variable, but a biological one, a property of the brain." [Jen99, p. 1] In this research, intelligence is approached from an evolutionary and biologically inspired perspective.

There is some general agreement that evolution is an important driving force of intelligence. In animals (specifically chordates), neural networks have served as the substrate for evolution of higher behavior. This is addressed in further detail in section 2.2.

1.3.2 Artificial Intelligence

The initial goal of Artificial Intelligence (AI) was recreating humanoid intelligence, but the majority of AI research has prescribed ideal mathematical

[7] I suggest that it is alarming. Let it be clear that by no means it is my intention to ridicule attempts at properly defining intelligence.

[8] The g factor is the highest-order common factor that can be extracted in a hierarchical factor analysis from a large battery of diverse tests of various cognitive abilities [Jen99], suggested by Charles Spearman in 1927 and still relevant.

behavior into computers [Swe03, p. 1].

Russell and Norvig clearly outline what the field of AI conveys. The definitions of AI vary along two main dimensions: first it is concerned with thought processes and reasoning, second it deals with behavior. Additionally, it is concerned with human performance or rationality (an ideal concept of intelligence). Table 1.1 shows the main four categories. [RN95]

The sub areas of AI range from fundamentals, like knowledge representation, knowledge acquisition, problem solving and search, to specific concepts, like knowledge-based systems, intelligent agents, natural language processing, machine learning, computer vision, or impacts. [RN95]

Table 1.1: Main A.I. categories

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

AI can be seen as having theoretical (science) and an applied (engineering) aspect. Most realized AI project include (1) A *theory* of intelligence, (2) a formal *model* and (3) a computer *system* implementing the model. [Wan06, p. 4]

Level of representation: connectionist versus symbolic AI

The best way to approach AI challenges is often a topic of vigorous discussion. The two main camps are connectionist[9] and symbolic AI. This is perhaps best illustrated by an example of speech recognition. For symbolic AI, form and meaning of words are represented in terms of labelled symbolic nodes [RMK06, p. 4]. Connectionists see the mapping of form to content as something that should develop in a neural network through training by examples. Connectionist pursue a bottom-up approach, while symbolic AI tends to analyze the higher level relationships and model those.

1.3.3 Motivation

The study of intelligence and AI has cross-fertilized many fields, including robotics, computer engineering (expert systems, conceptualization of the semantic web), neurology[10], psychology and even philosophy. Practical applications are OCR[11], speech recognition, autonomous mobile robots,

[9] Bottom up. Often also referred to as Parallel Distributed Processing.

[10] A notable discovery is e.g. [SD02]: measurements on a cat's visual cortex were fed into an ANN simulation. The simulation approximated the real world situation and served as a predictor for neuron behavior.

[11] Optical Character Recognition: recognition of hand- or typewritten text into machine-editable text.

driverless vehicles, etc. Accuracy and reliability of the former recognition systems is often still low (an order of a magnitude below human performance for speech [Lip97]). The fact that the later vehicles are often rated based on the 'mean time to intervention' shows that work needs to be done to make them truly autonomous and still able to deal with varying environments and incomplete information. Non-AI applications usually completely lack these properties, while they could be highly advantageous.

Through both highly theoretical (and often philosophical) and engineering work, we can gain insights into the very essence of ourselves and what is yet to come. While this constitutes a very long journey, I'm personally eager to take part in it and to be fascinated and perhaps surprised by it, just like nature is inspirational and surprising. I have a fascination for understanding 'the way things work' and, if possible, creating something new.

Chapter 2

Biology's driving principles for intelligence

There are many processes at many levels that contribute to the phenomenon called intelligence. These processes are approached from fields ranging from cellular chemistry and neurology to psychology and sociology. Instead of a complete treatment of these specific processes, I will focus on general principles that seem to appear throughout these processes.

Unmistakably, the brain plays a role that is central to intelligence. Instead of only investigating the biological brain as it is now (section 2.1.1), it also is interesting to know how and why it evolved into its current form (section 2.2). These evolutionary principles are critical to the subsequent chapter concerning artificial implementations. The observed increase of complexity and brain size are addressed, in sections 2.3 and 2.4 respectively.

2.1 Substrates

2.1.1 The Brain

Many researchers agree that the brain is poorly understood [Got97, p. 14]. However, numerous studies have been conducted and proved very valuable; more knowledge is acquired by the day. Luckily, complex and often complicated phenomena are found to have underlying principles that are not very complicated by themselves. The brain can be studied from an anatomical or a functional point of view. Anatomically, it must be pointed out that the neo-cortex is the latest evolutionary addition. It is only present in mammals. The cortex in total makes up for more than 80% of the brain's volume. It is the outer layer of the (fore)brain. It is 'folded' into gyri and sulci (bumps and cracks) to fit the area of 0.25 m^2 (equivalent of four A4

sheets of paper[1]) into our skull. [CGC+98, KW01, p. 56]. The cortex is essential to our higher functioning, because of its high connectedness.

It is an important principle that evolution adds layers and retains the old ones. Our neocortex is the latest addition. For coordinated movement, each layer adds a refinement to the more primitive ones. The more primitive layers provide reflexes and allows you to maintain posture without conscious intervention. For thoughts, each layer allows a higher level of conceptualization. The neocortex consists of an extremely dense set of neurons (circa 10^4 neurons/mm^3) [RMK06, p. 100]. The brain is organized for information to travel multiple paths even in a specialized subsystem, and lateral connections complement the received signals. Excitation and inhibition are the basic interaction of the 100,000,000,000 neurons[2]. Together with an even larger number of neuroglia[3] that support the neurons and help direct their growth.

Since our goal is to distil important properties of the brain, we will not discuss neuroanatomy any further. The function of brain and cognition is to enable the organism to attend to, process, and behaviorally respond to the forms of information and conditions that co-varied with survival or reproductive prospects during the species' evolutionary history [Gea05, p. 125].

An important feature that is often forgotten is that the brain is connected to a very versatile and capable body. It allows us to perceive a lot of our environment, and manipulate objects. The brain itself did not evolve by itself, it did so together with this body and, unique to humans, with complex tool use.

Perhaps the most prominent contribution to the understanding of neural dynamics was that of Donald Hebb [RMK06]. In commonly called 'Hebbian learning' a synapse between two neurons is strengthened when both the presynaptic (input) and postsynaptic (output) neurons are firing simultaneously. This form of self-organization is, despite its simplicity, now considered a general principle [Nol01, p. 51]. It allows associations of stimuli to arise, and serves to explain phenomena such as conditioning and associative learning[4].

2.1.2 The neuron and neuronal structures

The building block of the nervous systems is the neuron (or nerve cell), depicted in figure 2.1. It receives stimuli or 'input' (e.g. from another nerve

[1] ISO paper standard defines: A$n = 2^{-n}$ square meter, A4 $= 2^{-4}$ m^2, so 4*A4 = 4*$^1/_{16}$ = 0.25 m^2

[2] [BK92, p. 7] say 10^{11}, [KW01, pp. 47, 79] say 80 billion neurons (order of 10^{11}).

[3] The earlier neuron count excludes another 10^{11}, perhaps even 10^{12} neuroglia (see also: http://www.sfn.org/index.cfm?pagename=brainBriefings_astrocytes).

[4] Memory formation is a more complicated subject.

cell) at its dendrites. Inputs can be exitatory (additive) or inhibitory. Once a certain threshold is reached, the neuron will propagate an action potential through its axon. Action potentials speeds range between 0.55 m/s and 111 m/s[5]. The myelin sheets cover the axon to minimize the amount of energy needed (and corresponding ion influx and efflux) to propagate the action potential at high speeds. Finally, the action potential will release neurotransmitters at its axon terminals which may be connected to subsequent neurons or to the muscles.

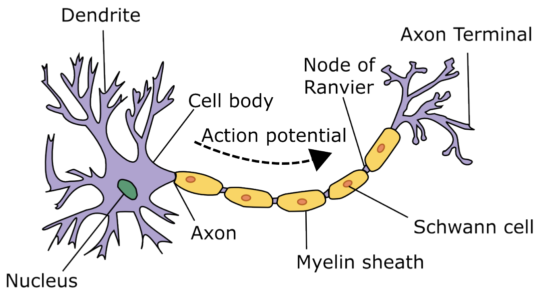

Figure 2.1: A neuron covered in myelin, conducting an action potential in the rightward direction. Adapted from: Dhp1080 on Wikimedia (originally adapted from: 'Anatomy and Physiology' by the US National Cancer Institute's Surveillance, Epidemiology and End Results (SEER) Program).

Also, the interconnectedness is very high; in the order of 10^4 synapses on average per neuron [BK92, p. 7]. There are various types of neurons, but they respect the general features mentioned above. The cells in the neuronal population have varying types of "transistor" properties. Cells in the cerebellum are fine-tuned and worked out for motor effectors, cortical cells are pyramidal and have their own characteristics and so do 'association neurons' in the thalamus.

2.1.3 Substrates in general

"To find real complexity on the scale dimension, we may look at the human body: if we zoom in we encounter complex structures at least at the levels of complete organism, organs, tissues, cells, organelles, polymers, monomers, atoms, nucleons, and elementary particles.". [Hey96]

I could dedicate an entire chapter to the intricacies and suitability of various structures or 'substrates', but this is not necessary for our purpose. It suffices to say that at their level, each substrate is very important - if not critical - to what originates from neuronal structures, the neuron, the

[5] 111 m/s [EA03, p. 1]. Perhaps up to 120 m/s [KW01, p. 132]. 100 m/s [Ham03, p. 61].

biological cell, proteins or genetic substrates such as DNA and RNA. Effects of low level impairments can cascade and expand over time [Elm05, p. 115]. Thus, for each of these structural levels there is an evolutionary pressure to improve internal organization and interactions with other levels. The internal organic medium is a basis for evolution and so is the external biotope. The structural levels extend up to the macroscopic ecology, providing a wide range of substrates spanning 12 orders of magnitude of scale at which evolution simultaneously does its work [Ray99].

2.2 Evolution

"If you consider all the biochemical steps required to get a message across a synapse, [...] you may wonder why such a complex communication system ever developed. The answer must be that this arrangement makes up for its complexity by allowing the nervous system to be flexible about the behavior it produces."

`Bryan Kolb & Ian Wishaw` (2001) [KW01, p. 162]

The human evolution highlights the importance of flexibility mentioned by Kolb and Wishaw (see also section 2.3). But what is evolution precisely? Evolution is a process that is applicable to biology, culture and even elementary particles. Unless stated otherwise, by evolution, we mean biological evolution theory[6]. While this disambiguation may seem unnecessary, evolutionary processes happen at various levels of scale, *even in biology.* For example, processes of learning encompass "the long internal process which changes structural properties of its carrier system"[Wik07a]. Perhaps it is not so intuitive, but this is certainly applicable to development, learning and memory-formation in the brain. To avoid confusion with biological evolution in the regular sense, we will call these Darwinian processes. Calvin describes a Darwinian process as: "A pattern (spatiotemporal firing pattern of a Hebbian cell-assembly, in this case) that copies with occasional variation, where populations of the variants compete for a limited work space, their relative success biased by a multi-faceted environment (both memorized and real-time, in this case), and with further variations centered on the more successful of the current generation (Darwin's inheritance principle)." [Cal98] The brain allows patterns and ideas to compete across lateral connections. Multiple, perhaps conflicting or incomplete signals are often interpreted without much effort. We are wired to perceive and learn from our environment, to be able to adapt to it.

Evolution, or 'descent with modification' [Cha01], is thought to consist of several elements [Cal04, Ch. 10]:

[6] Defined as: "Biological evolution is the process of change over time in the heritable characteristics, or traits, of a population of organisms." – source: Wikipedia

1. A pattern (or substrate in our terminology)
2. The pattern (or substrate) is copied.
3. Variations occur, typically from copying errors and recombinations. Some of this variance is heritable [Cha01].
4. Variants compete with eachother for finite resources or space.
5. Finally, the next round (generation) is centered around those variations that proved to be more successful (at reproducing, not just at surviving).

Element 3 is non-deterministic, or 'random'. The final element is not. Furthermore, opening (biological) niches due, climate change, sexual recombination, speciation[7] and temporary inbreeding (island populations) catalyze the evolutionary process [Cal98, Cal04, ch. 10].

2.3 Evolution and complexity

"Nature works by steps. The atoms form molecules, the molecules form bases, the bases direct the formation of amino acids, the amino acids form proteins, and proteins work in cells. The cells make up first of all the simple animals, and then the sophisticated ones, climbing step by step. The stable units that compose one level or stratum are the raw material for random encounters which produce higher configurations, some of which will chance to be stable... Evolution is the climbing of a ladder from simple to complex by steps, each of which is stable in itself."

`Jacob Bronowski (1973)`

Also sharing this view on stratum theory and its implications are, Butler (1878, 1880), van den Tweel (1988), Heylighen [Hey96] and Vroon [Vro89, p. 175]. It corresponds with the observation that cognitively higher species (which exhibit more intelligent behavior) tend to be more recent products of evolutions [MSS95].

Does evolution lead to greater complexity? Since it *is* a general tendency, it is often erroneously believed to always be the case. In some cases, simplification allows for more robustness, efficiency and specialization, both in a biological and digital evolutionary medium [Ray94, p. 14]. But when evolutionary pressues demand an adequate response to variant circumstances, complexity is favored [Gea05].

[7] The evolutionary formation of new biological species, usually by the division of a single species into two or more genetically distinct ones. Source: The American Heritage Dictionary of the English Language, Fourth Edition

Functional complexification follows from the need to increase the variety of actions in order to cope with more diverse environmental perturbations, and the need to integrate actions into higher-order complexes in order to minimize the difficulty of decision-making [Hey96, Gea05]. Both processes produce a hierarchy of nested supersystems or metasystems, and tend to be self-reinforcing. Though simplicity is a selective factor, it does not tend to arrest or reverse overall complexification. Increase in the absolute components of fitness, which is associated with complexification, defines a preferred direction for evolution, although the process remains wholly unpredictable. [Hey96]

2.4 Intelligence and brain size

There was nearly a threefold increase in brain size from apes to modern humans [KW01, p. 24], roughly according to an exponential growth trend[8]. There are several things that we do better than apes, that are demanding for the brain. These differences may provide insights into our intellectual growth. Performing well in complex, diverse and most notably novel circumstances requires generally useful cognitive faculties[9]. Problem-specific skills are not as useful when facing a new challenge. Static or invariant conditions create pressures for the evolution of modularized systems[10] whereas dynamic conditions create pressures for modular plasticity and the evolution of less modularized, domain-general systems [Gea05, p. 125].

A recent meta-study by McDaniel confirms a significantly positive relation (0.33) between brain volume and (general) intelligence [McD05]. Azam nuances this by saying that "[i]t can be deduced that an increase in brain size does not necessarily increase the sophistication or behavioral diversity, unless accompanied by a corresponding increase in specialized brain modules" [Aza00, p. 15]. Others suggest that the correlation of brain volume to intelligence is poor and that factors such as glia-to-neuron ratios[11] can serve as a possible determinant [KW01, p. 572]. There is a moderate correlation between IQ and kinship[12] ($r = 0.86$) for identical twins. This influence increases with age [Sha02, pp. 89-92]. This shows that intelligence is heritable and, when useful, evolution will favor the more intelligent.

Macroscopically, humans were quite flexible and often filled biological niches that became available. Purely (evolutionary) biological changes sug-

[8] Based on archaeological skull research. [CGC$^+$98]

[9] Development of beliefs and mental models of the world, generalization, association, causality, logic, reasoning.

[10] Modularized systems are the more specialized systems, as opposed to domain-general systems.

[11] Glia are cells that provide nutrients to the neurons. They are in-between the blood veins and the neurons.

[12] Kinship: Measure of genetic relationship.

gestions have also been made. The increase of brain size is constrained by a problem of heat dissipation. The so called 'radiator hypothesis' states that more adequate cooling that can be found in humans may have made this increase possible [Fal05].

Based on recent datasets (contemporary subjects), larger brain volumes were associated with higher general fluid intelligence[13] (r = 0.49), larger short-term memory capacity (r = 0.45), and faster speed of processing (rs about -0.4) but were unrelated to general crystallized intelligence[14] (r = 0.06) [WVL00, Gea05, p. 114]. Fluid intelligence related traits seems to have arisen together with the increase in brain size.

Many suggest language is important to the development of intelligence [Cal04, KW01, p. 533], others suggest it is not critical [Pen90, p. 414]. The level of toolmaking seems to be an indicator of intelligence, but not a driving force. During periods of conservatism, without advances in toolmaking, brain size had been gradually increasing, so it appears it did not significantly stimulate an increase in brain size. The level of tool *usage*, however, differs dramatically between early and later hominids. Throwing objects requires planning a motion ahead and perfecting it under different distances, with differing object weights, etc. In hunting, improvements were directly rewarded with food [Cal04, Ch. 8]. Taking from the catch of others is often accepted if one contributes. Keeping track of this required improvements in memory and language.

The precise reasons for intelligence can not easily be derived from a physiological view. Some evolutionary events have lead to plausible explanations for the development of our intelligence.

2.5 Development

In contrast with the extreme view of a newborn's state being a *tabula rasa* or 'clean slate'[15], evidence suggests that humans are equipped with several highly evolved preconditions for abilities from the moment of conception[16] (based on the genotype). From the moment of conception, interactions with the environment develop these features, shaping the phenotype. The sensory modalities vision, hearing, touch and smell provide stimuli. The stimuli direct the wiring of neurons, and areas of the brain are recruited that become specialized in precessing this information. Interconnection of these specialized systems or 'modules' is highly plastic during early development.

[13]General fluid intelligence: the ability to solve problems and to learn.

[14]General crystallized intelligence: Acquired knowledge of the world, (e.g. commonsense, mental models, vocabulary, etc.).

[15]In AI, tabula rasa is also used to highlight that an AI system is not programmed with facts but arrives at its conclusions by its internal dynamics.

[16]Conception is the initial stage [Sha02, pp. 102], not birth recognizing prenatal influences.

The experiences needed to adjust these plastic features to these ecologies are generated by children's natural social, play, and exploratory activities [Gea05, p. 125].

Learning in humans occurs on-line, as opposed to batch learning used mostly in machine learning. The balancing of learning (exploring) and using what is learned (exploiting) is very important to development. In section 3.2.2 a simulation of 'infant development' will be given. On-line learning is important in this example.

Chapter 3

Artificial Implementation

Human-like intelligence relies on machinery of high sophistication in terms of processing capabilities. In this chapter I'm going to discuss artificial counterparts to two mechanisms that have played a key role in the development of intelligence in biological systems: evolution and neural networks (section 3.1).

In an attempt to understand and a hope to eventually harness the strengths of this machinery, artificial neural networks and evolutionary algorithms have been created. Different versions of these mechanisms have been created to answer different questions or to harness different qualities of these algorithms.

Neural networks have been modeled at the level of the synapse to answer neurobiological and neurophysiological questions and to test hypotheses about the workings of the brain. Other implementations were used to classify images or patterns. The most common implementation of the artificial neural network (the canonical neural network), however, is not the most detailed nor biologically plausible version. The canonical neural network will first be presented, and the discussion that follows will address whether essential features are left out, or whether other implementations are feasible.

Implementation of another instance of evolution, next to natural evolution, brings forward many questions [Ray99]. What assumptions can be made, on what basis can we compare it, etc. Again, multiple perspectives on evolution result in different implementations of artificial evolution (see also section 4.2). Like with neural networks, a common evolutionary algorithm will be used as a reference example.

3.1 Artificial counterparts of key mechanisms

3.1.1 Evolutionary algorithms

Evolutionary algorithms (EAs) are a biologically inspired set of rules that describe the use of an evolutionary process in computation. In the typical

EA, during a generation, members of a population are ranked according to a fitness function. Those members with the highest fitness ranking are given a higher chance to become parents for the next generation, the offspring. The exact method used to generate offspring from the parents, is termed the reproduction heuristic. The common reproduction heuristic is a mutation rate (per chromosome) and sometimes also a chance of having mutations. After a number of iterations, a typical population increases its fitness and converges towards (local) optima. EAs function very well at optimization when the search space is large.

3.1.2 Genetic algorithms

Genetic Algorithms (GAs), as developed by John Holland [Hol75], are Evolutionary Algorithms that use genetic recombination as the main reproduction heuristic, accompanied with mutation.

With GAs, biological consistency is maintained to a further extent than EAs, especially when n-point crossover is used. Uniform Crossover, a given probability of cross-over at any gene, is not biologically consistent and is not widely used. There are also parallel reproductive heuristics, where multiple populations are evolved but mostly (or entirely) kept separate, Calvin [Cal04, Ch. 10] argues this has been an important extra catalyst to natural evolution. For an overview of the different reproductive heuristics read Gordon and Whitley [GW93]. All known species in nature either produce asexual (one parent) or sexual (two-parent) recombination. This restriction is not a requirement in artificial algorithms and improvements in performance have been reported when increasing the number of parents [ERR94]. Sex is another catalyst, but not an essential according to Calvin [Cal04]. Box 3.1.1 displays the pseudo-code of Holland's Simple Genetic Algorithm (SGA). [MT95]

Algorithm 3.1.1: GENETICEVOLUTION(P)

$t \leftarrow 0$;
$initialize(P(t=0))$;
$evaluate(P(t=0))$;
while $isNotTerminated()$

$$\textbf{do} \begin{cases} P_p(t) \leftarrow P(t).selectParents(); \\ P_c(t) \leftarrow reproduction(P_p); \\ mutate(P_c(t)); \\ evaluate(P_c(t)); \\ P(t+1) \leftarrow buildNextGenerationFrom(P_c(t), P(t)); \\ t \leftarrow t+1; \end{cases}$$

end

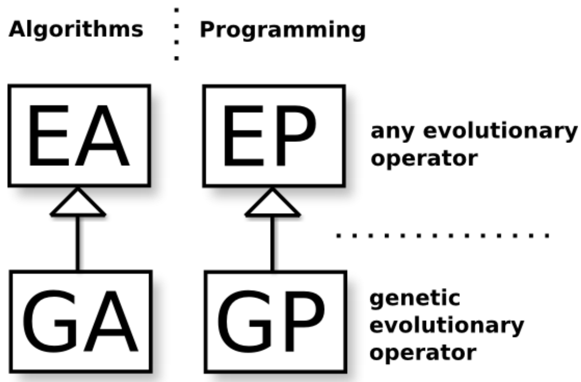

Figure 3.1: Evolutionary Algorithms (above left), their subtype Genetic Algorithms (bottom left), Evolutionary Programming (top right) and Genetic Programming (bottom right).

Evolutionary algorithms are also used to directly evolve lines of computer programming code, this is called Evolutionary Programming (EP) and Genetic Programming (GP) when recombination is used. The relation between these four domains is displayed in UML notation in figure 3.1.

In section 3.2.1 I will present an example of an application of a genetic algorithm.

3.1.3 Artificial Neural Networks

Artificial Neural Networks (ANNs) have multiple representational forms. The most common are the mathematical (equation 3.2) and graphical form (figure 3.2). For *each* Artificial Neuron, the mathematical form consists of a function $g(x)$ of the input vector x, where $x = (x_1, x_2, \ldots, x_i)$. Each input x_i is weighted according to its weight $w = (w_1, w_2, \ldots, w_i)$. K is the post-processing function that is finally applied. This results in the following equation for a single neuron:

$$g(x) = K\left(\sum_i w_i x_i\right) \tag{3.1}$$

Neural Networks consist of multiple artificial neurons like these. The output of one neuron is connected to the input of another neuron. Mathematically, $g(x)$ (the result of equation 3.1) is an input to a neuron with function$f(x)$:

$$f(x) = K\left(\sum_i w_i g_i(x)\right) \tag{3.2}$$

To introduce non-linearity, a hyperbolic tangent or sigmoid (S-shaped) function is commonly used for K. Non-linearity is deliberately analogous to

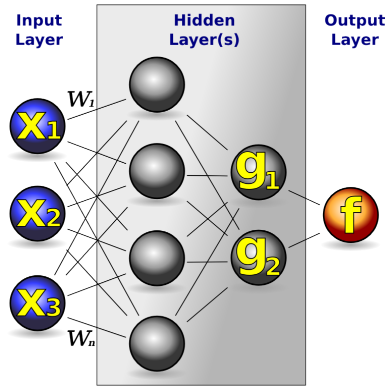

Figure 3.2: An example graphical representation of a Multi-layer Feed Forward Neural Network.

biological neurons, and responsible for its versatile information processing properties.

A reason for its popularity [Wik07b] in neural networks is because the sigmoid function satisfies the differential equation $y' = y(1 - y)$. The right hand side is a low order polynomial. Furthermore, the polynomial has factors y and $1 - y$, both of which are simple to compute. Given $y = \text{sig}(t)$ at a particular t, the derivative of the sigmoid function at that t can be obtained by multiplying the two factors together. These relationships result in simplified implementations of artificial neural networks with artificial neurons.

3.1.4 Evolving neural networks (EA and ANN combined)

Intelligence cannot be attributed to a distinct concept such as EA or ANN alone. Similar to the vigorous nature versus nurture debate, both can be attributed to play an important role, while it is the combination that results in the end product.

Evolution is agnostic about what to develop - it has no design goals [Hop82, p. 2254] - but qualities do emerge that have co-varied with the prospects of persistence of this quality. The prospect of persistence is normally related to the reproductive and survival prospects of the individuals in a population [Gea05, p. 99]. In artificial systems it is determined by the fitness function. In section 4.2, possible issues with the fitness function are further discussed.

Evolutionary algorithms are often employed to solve optimization problems which are otherwise computationally prohibitive/expensive (np-hard and np-complete problems). While they are useful in this respect, it greatly undervalues their potential of being able to develop solutions to any challenge. Evolutionary algorithms have no difficulty evolving circuits beyond

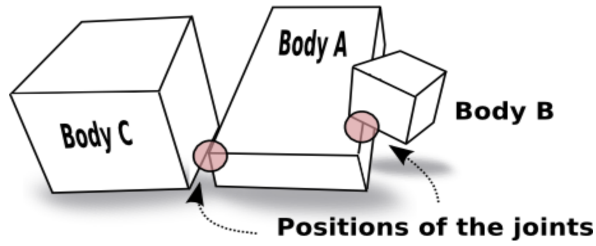

Figure 3.3: The gentype of the virtual creature. Adapted from [GG07] with permission.

the complexity that human designers can grasp [Hop82, p. 2256]. It is critical to note that evolution and artificial evolution succeed in finding a solution to problems without explicitly giving it knowledge on how to solve the problem and regardless of what the problem is. I argue that this generality is important to intelligence and to achieving artificial intelligence.

3.2 Example implementations

3.2.1 Genetic evolution of virtual creature morphology and control

"Instinct is intelligence incapable of self-consciousness."
John Sterling

Goldstein and Godoy [GG07] have evolved creatures of varying shapes[1] and joint positions and strengths. This first example will illustrate the ability of genetic algorithms to differentiate and optimize the physiology of virtual creatures.

The implementation

A genetic algorithm evolves a population of virtual creatures. At first it is randomly initialized (with some constraints such as maximum number of body-parts). Each genotype (genetic code) prescribes exactly the characteristics of the individual, the phenotype. Besides body-part sizes, control of the artificial muscles connected to the joints is also determined by the genotype (see figure 3.3 and figure 3.4). The muscles are controlled by a sensomotoric neural network, sensing and actuating muscles. The controlling "nervous system" is non-central, each joint has its own independent neural network.

[1] Morphology: concerning the form or shape in general, or physiology in biology.

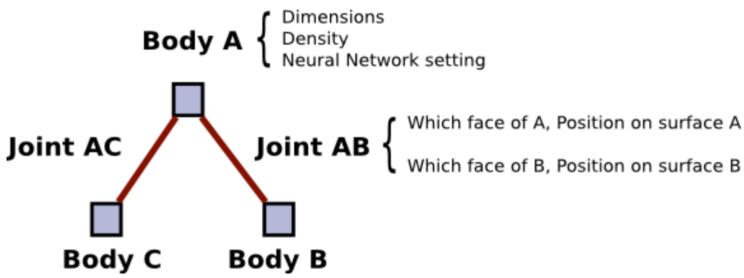

Figure 3.4: The phenotype of the virtual creature. Adapted from [GG07] with permission.

The selection process is based on the distance traveled from the starting point. This distance is evaluated by a simulation in a physics engine. The calculations are normally non-visual, but can be visualized in 3D by a tool (see figure 3.5)[2].

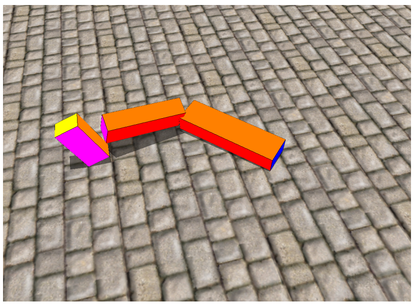

Figure 3.5: A screenshot of the EvoMorph visualization tool rendering a crawling creature in 3D with openGL^{TM}.

Lessons learned

It is important to realize that there are many implementation issues, that one may not think of at first. How do you measure distance traveled? Do you measure from the center of the body or the outside? If the creature is larger or very long, this scale advantage may not be fair. Using the wrong Cartesian coordinate will result in creatures with jumping behavior as their main trait, instead of locomotion. Is distance based on the final position

[2] A short video clip is available at: http://studie.erikdebruijn.nl/thesis/

vector minus the initial position vector, or do you sample in-between? Creatures that are very fast but run in circles would lose from slow but straight movers. Sampling positions very often might award a 'trembling' creature with movement, while it stays on the spot in longer timespans.

The previous illustrates that there are many considerations when designing a selection function. You get what you select (very literally), but the way in which the creatures do it is creative and may be unexpected. Not only is the genetic algorithm able to evolve better locomotion strategies, it is also able to "find" and exploit inaccuracies in the physics engine to achieve a high fitness. It is very important to think through the fitness function. This may include formulating a penalty for cheating the physics engine (causing unrealistic simulated behavior). Especially with repetitive motion patterns rounding errors can accumulate [KN93].

It also matters a great deal what reproductive heuristic you choose. Do you recombine only properties with the top-ranking? If so, the whole population might become quite homogeneous after a while and only a (very) select part of the search space will be searched (very) thoroughly. In other words: a lot of 'strategies' will not be 'tried', there will be at most fine-tuning of a main theme (which is itself quite random). Would you recombine two parents or a different number?[3] Another important observation is that after some generations, the creatures do not immediately converge into only one shape with one locomotion strategy. For a long time various strategies will exist side by side. Many 'strategies'[4] arise, like hopping, walking, rolling, pushing, swimming.

Relation to other work

The relation to the evolution of biological species is clear, however it is highly simplified. There is no interaction between creatures and goals are static. No actual learning occurs during a lifetime: only changes to the input of neural networks, no changes to the structure. However, knowledge *is* accumulated across generations.

There are many examples of research into Artificial Life[5]. Thomas Ray's Tierra shows an ecology of programs competing for resources. Parasite-host,

[3] An apt quote in this regard is: *"Sex is a relatively recent addition to the dance of life. For more than 2,000,000,000 years, asexual reproduction was the rule. You know, if you were a creature, you just separated into two clones."* – `Mark Jerome Walters`

[4] When I say 'strategies' I do not conjecture that they have thought of a way to reach a goal. The experimenter has a goal, but the algorithms are just procedures which should not be ascribed any motivations. Creations are often ascribe more human-like properties than can scientifically be justified [SZ98].

[5] It is disputable whether the example may be labeled Artificial Life, since there is no adaptation, socialization nor is there growth. Oparin [Opa38] proposed that living *matter* be defined as having the following properties: metabolism, self-reproduction, and mutability.

parasite-hyperparasite relationships arise, etc. Ray's Tierra is an example of 'programs' that evolve native to the digital medium. The former example illustrates, by analogy, the evolution of insect-like[6] creatures. While this example makes use of both ANNs and EAs, all behavior is purely instinctive. While some individuals may appear to act intelligently, this is certainly not consciously. Of any individual it would depend on chance if it would cope well under novel circumstances.

3.2.2 Babybot: an artificial developing robotic agent

> *"[...] It can also be maintained that it is best to provide the machine with the best sense organs that money can buy, and then teach it to understand and speak English. This process could follow the normal teaching of a child. [...]"*
> `Alan Turing (1912-1954)` [Tur50]

Metta et. al. [MPMS00a, MPMS00b] describe their experimental robot which is displayed in figure 3.6. The theoretical foundation of Metta e.a. are rooted in the 'Piagetian' theory of Jean Piaget, psychologist with a biology background. He is widely known for his theories on the development and 'the teaching of a child' (in 'The origin of intelligence in children' [Pia36]). He defined intelligence as a basic life process that helps an organism to adapt to its environment [Sha02, p. 50].

The example of evolving morphology and control is biologically similar to insects with more or less deterministic, static responses to stimuli. Evolution produces a 1 on 1 mapping of genotype to phenotype. To model higher level organisms such as a humanoid baby, with its complex brain structure, it is essential to take into account the interactions of development of the phenotype of the individual. This 'development' allows the emergence of qualities such as intelligence.

The implementation

Another difference is that Babybot is physically realized as a robot. This essentially means that the "physics are free". This nullifies the problems of inaccuracy of simulated physics and perhaps an overly homogeneous or unrealistic environment (such as in the virtual creatures example).

It was constructed to be a test bed for various theories of human development, and is thus constructed with biologic realism in mind. For example, the robot observes the world through a high-resolution fovea and a progressively lower resolution periphery. Furthermore, it includes two microphones,

[6] In terms of the level of simplicity, not its specific morphology, which is always six legged for insects.

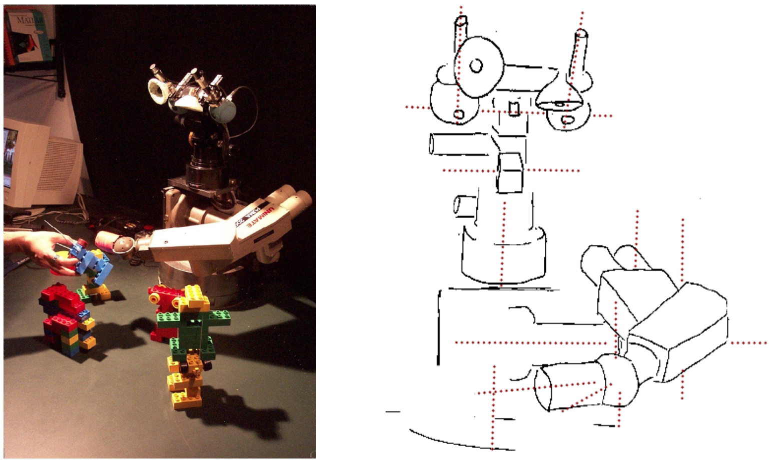

(a) (b)

Figure 3.6: The LiraLab Babybot. (a) Babybot interacting with its environment. (b) Schematic version showing its degrees of freedom (DoF)

an inertial device and motor encoders (that can effectuate movement for 12 DoF[7]).

Lessons learned and relation to other work

Even if it were possible for a genetic code to convey the complexity of the end product, which is isn't [Cha96, pp. 9, 15], it would be highly inefficient. Nature solves this by encoding growth patterns, not the end product. Boers and Kuiper [BK92, Chapter 4] emphasize that the recipe is genetically stored, not the blueprint (also [Cha96, pp. 9, 15, 16]). They model this recipe in terms of L-systems which can approximate growth[8]. It is well known that modularization has occurred in many biological systems[9]. The brain is not an exception. Modularization plays an essential role, for stability and also for efficiency (modules can be repeated) [Aza00]. The brain can be considered modular as well. The importance of a 'bootstrapping process' is highlighted by Metta et.al. The modules are plastic in critical periods and become more stable over time. The development of modules later on is highly affected by earlier modules. [MPMS00b] "[Newborns] show a series of 'innate' behaviors, basic control synergies and reflexes[10]." [MPMS00b].

[7] In mechanics, degrees of freedom (DoF) are the set of independent displacements that specify completely the displaced or deformed position of the body or system.
Source: Wikipedia (retrieved 15 October 2007).

[8] Lindenmayer-systems are introduced by Astrid Lindenmayer in an attempt to model the biological growth of plants.

[9] Think of organs, cells, chromosomes, etc.

[10] A reflex is an involuntary and automatic response to a stimulus, such as when the eye automatically blinks in response to a puff of air. [Sha02, p. 134]

The primitive reflexes are controlled by the lower 'subcortical' areas of the brain and are lost[11] once the higher centers of the cerebral cortex mature and begin to guide voluntary behaviors [Sha02, p. 137]. The robot explores and exploits its environment simultaneously. Balancing the trade-off between exploration and exploitation is an important problem [KLM96, p. 243]. When eye movements become reliable and consistent the neck started moving as well, which provided feedback from the inertial sensors. The goal of the designer has shifted to devising a suitable initial state (at time t_0), and the appropriate developmental rules to get some close approximation of the desired 'final product' as opposed to building the final product itself [MPMS00b]. The robot faced problems of overshooting targets and corrective oscillations. The time slots for enabling modules were explicitly but carefully programmed [MPMS00a, p. 9].

This research provides valuable insights into the process of learning of (spatial) perception and originating abilities. For further work it would be interesting to know whether the timing of this 'bootstrap process' is genetically orchestrated (literally prescribed) or whether the timings purely depends on the reliability and consistency of a developing module (a dynamic and variant property of individuals). The above experiment[12] does not yet validate whether this is the case, but makes it plausible. Moreover, neuronal growth mechanisms are 'challenged' to create a representation (exploration) of a real-world problem domain, while concurrently exploiting it (making use of what is learned). "As we examine, [development] is a uniquely powerful and general learning strategy that undermines the central assumptions of classical learnability theory, which is premised on the assumption that the learning properties of a system can be deduced from a fixed computational architecture." [QS82, p. 23]

3.3 Barriers to implementation

3.3.1 A suitable substrate

Implementation of plasticity

No one will argue against the fact that the 'software' in every brain is different, because it depends on the actual experiences to take shape. However the 'software' is not stored in special purpose memory gates (that can contain just any binary encoded value). The behavior that originates from the brain is 'encoded' in its structure:

[11] Or, [MPMS00a, p. 9] suggest they may become embedded into more complex control structures.

[12] Similar to the my own (unfinished) experiment 'RoboSense Experiment 1' on http://studie.erikdebruijn.nl/thesis/

"Learning does not just alter the knowledge base for a fixed computational engine,
it alters the internal computational architecture itself."
`Andy Clark` [Cla01]

Plasticity (in the sense of changing interconnection topology) is fundamental to learning in biology.

Currently, when creating a processor device, all attempts are made to make sure every piece of hardware is the same. Elimination of variation is a central thought when manufacturing computer chips. A microprocessor with too much variability[13] in its substrate, would render it useless. By design it depends on deterministic behavior of its components. Biological neural structures, which are plastic, are much more robust and can even make use of the physical diversity of its micro-environment [CdZR93]. An enormous contribution of von Neumann to computing is to store data and executable code both in memory, lifting the limitation of application specificness of computers[14]. While operations and data are on the same medium, reliable execution demands that data is never mistaken for operations or vice versa. This has been an important source of problems and exploitability in software. The brain doesn't work based on a clear distinction. To the brain, memory and experience is inherently meaningful (which data by definition is not) and greatly determine the way we interpret and perceive. Interpretation and perception of the brain can hardly be considered programs performing operations. This illustrates a fundamental difference of the architectures of the traditional computer and the brain. A computer's memory resides in distinct units, while the brain encodes every memory in its structure.

The former is a comparison of traditional computing to the brain. Obviously, simulating neural behavior (with ANNs) would be more similar to the functioning of the brain. But the logical and physical separation of memory units and processing units puts considerable requirements on the I/O speed and bandwidth of traditional architectures. A possible way out of this is using a different VLSI chip architecture: FPGA[15]. An FPGA is essentially an array of memory that manifests as logic gates. Next to the logic areas, there are (re)programmable interconnect blocks that constitute the switching areas. The discrete state nature of digital computers has been considered a problem. Analog VLSI (aVLSI) is another approach which

[13] e.g. due to impure silicon, lithographic mask alignment variations, a scratch on a lens or dust during production.

[14] "In a special purpose machine the computational procedure could be part of the hardware. In a general purpose one the instructions must be as changeable as the numbers they acted upon. Therefore, why not encode the instructions into numeric form and store instructions and data in the same memory? This frequently is viewed as the principal contribution provided by von Neumann's insight into the nature of what a computer should be." [Ril87]

[15] FPGA is the acronym for **F**ield-**p**rogrammable **g**ate **a**rray

does not have the problem of computationally expensive multiplication by logic gates [ZS03].

Neural Networks are characterized by parallelism, modularity and dynamic adaptation. FPGAs are well-suited because of concurrency, and reconfigurability. The reconfigurability FPGA aspect is exploited in several ways. In a sense, neuroplasticity is achieved by topology adaptation. Learning is implemented by adaptation of weights. FPGAs also provide a good basis for rapid prototyping of ANN designs [ZS03].

Massive parallelism

Perhaps the most distinguishing feature of the brain in comparison to regular computing is massive parallelism. While primarily a quantitative difference, it has significant implications to the character of processing in the brain. This can be illustrated by the example of recognizing a figure. This process normally occurs so fast that given the speed of neural interconnections, only something in the order of a hundred consecutive processing steps could have taken place [Ham03, p. 11]. Meaningful integration of many concurrent processes is important to relatively slow, in comparison to electronic circuits, neural communication.

Whereas in conventional computers synchronization of the digital building blocks is achieved using a clock signal, there is no such global clock in biological systems. Processes can execute slowly and unconsciously simultaneous with others that are conscious. In a more biologically oriented simulation, global synchronization should be avoided [Roj96].

Currently, the maximum reported scale based on FPGA seems around 65 million simulated neurons[16]. The initial goal of 1 billion can still be reached, since it's not a constraint of the architecture [dGKG$^+$98, dGK02]. Importantly, such a larger number of neurons could be simulated on an ordinary computer, but it would very slow instead of real-time.

Rich connectedness

The brain is also praised for its ability to integrate many signals. This is not just an opinion because the average has between 10.000 to 100.000 synapses (human brain)[17]. The architecture of brain very different from serial processors. Other routes have been suggested and are pursued: Parallel computing, neuromorphic engineering, optic-holographic [PWB87][18] [Nol01, pp. 55, 189-195] and molecular [CdZR93]. Molecular computers may solve the scaling problem.

[16] As far as I could find. There are some false claims of 1 billion neurons made [Kur99, p. 80] and recited [Nol01, p. 54], but that was the *unrealized* goal.

[17] [BK92, p. 7] say: 10^4 synapses on average per neuron.

[18] The first optical implementation of neural networks was proposed by D. Psaltis, cited here.

Apart from a medium for 3D pictures, holography may be a promising method of optical interconnection.

“In contrast with implementations, in which the specifications of the connection patterns must be stored separately from the connections themselves, holographic media can simultaneously provide both the massive physical interconnectivity and the large memory required to specify the connections. This duality is particularly useful in adaptive networks.” [Ram98]

3.3.2 General intelligence

> *”How strange that our most advanced systems can compete with human specialists, yet be unable to do many things that seem easy to children.”*
>
> `Marvin Minsky`

There have been many achievements by symbolic approaches in AI. Most of these are focused on the category “systems that think rationally” (as discussed in section 1.3.2) and are highly specialized. “[M]ost well-structured problems such as textbook math and science problems, are simple because they tend to engage a constrained set of variables that behave in a predictable ways.” [Ese06]. Some ascribe intelligence to an agent capable of solving diverse problems [LH06]. This would apply to Simon and Newell’s General Problem Solver (GPS) and the universal proof searcher, the Gödel machine (see section 4.1.3).

There are two types of intelligence: fluid or crystallized. Roughly, the former is categorized as the ability to learn or to solve problems, the latter is that which is learned. General intelligence (Spearman’s g-factor) is highly related to the fluid type of intelligence. Moreover, there is a strong neural basis for this relation. The general form of intelligence is highly influenced specifically by a frontal system [JD00].

Chapter 4

Emergence of intelligence

This chapter attempts to address whether the former techniques of artificial implementation (presented in chapter 3) actually give rise to intelligence. Since no conclusive answer can be given within the scope of this thesis, I will bring forward two views on the possibilities and limitations of the artificial mind.

The computability of thought will be discussed on from the perspective of theoretical computer science, by introducing Gödel's theorem and from the point of view of philosophy by introducting John Searle's 'Chinese Room' thought-experiment (section 4.1). Can thought, consciousness and the mind be regarded as algorithmic? What are the implications if thought is non-computable? Finally, I will bring forward some ideas from Alastair Channon (section 4.2). He argues for using an alternative to the traditional fitness function in order to achieve intelligence.

4.1 The computability of thought

I propose to consider the question,
"Can Machines Think?"
`Alan Mathison Turing`

4.1.1 Intelligence and thought

There is an important difference between acting intelligently, and being intelligent. An agent acting intelligently can rely on execution of a prescribed or predetermined repertoire, similar to the 'seemingly intelligent' instincts of Goldstein and Godoy's evolved virtual creatures. Actually being intelligent requires a level of understanding originating from thought. In this thesis, we've identified algorithmic methods that may contribute to the emergence of intelligence. While some rudimentary intelligence (or intelligent instincts)

appears to arise, concepts such as 'thought' introduces a range of deep philosophical questions.

4.1.2 The Turing test and intelligence

For the assessment of intelligence Alan Turing has devised a formal test which is now widely known as the Turing Test (for a short explanation see Appendix A). In short: "It is proposed that a machine may be deemed intelligent, if it can act in such a manner that a human cannot distinguish the machine from another human merely by asking questions via a mechanical link." [Abe98] The 'intelligent machine' hypothesis initially gained momentum by the increase of knowledge on the biological brain. The neurobiological processes were first thought to be similar to or identical to the information processes of a computer, since the neuronal function resembled that of a logic gate, so a computer made of logic gates was allegedly capable of simulating an intelligent mind.

4.1.3 Formal proof and Gödel's theorem

Is the mind a product of logic gates? Investigating these deep questions entailed formalization of some key concepts. Concrete mathematics provided a foundation to this philosophical discussion, but also provide insights into potentially fundamental constraints.

There is an tough debate around the generation of proof of propositions. In some cases, an algorithmic machines seems uncapable of solving something that a human can.

In order to unequivocally be able to decide the truth of any proposition, Hilbert created the challenge of once and for all creating a sound and consistent formal approach. If successful, anything that was provable would be computable. Instead, the converse was proved by Kurt Gödel[1]. His incompleteness theorem showed that in any language expressive enough to describe the properties of the natural numbers, there are true statements that are undecidable: their truth cannot be established by any algorithm. Importantly, it is not a shortcoming of any particular formal system, but rather a property inherent in all formal systems [Haz02, on undecidability].

Many 'procedural' mechanist operations, such as multiplying numbers, can be executed by a computer faster, more precisely and more efficiently. Gödel's first and second theorem tell us that there is a thing that no algorithmic machine can do everything that the human mind is capable of[2]. Gödel did not deprive scientists of all hope to create an artificial mind, but he did

[1] Sometimes called "father of theoretical computer science". Source: http://en.wikipedia.org/wiki/Timeline_of_artificial_intelligence

[2] "[W]e can see that the Gödelian formula is true: any rational being could follow Gödel's argument, and convince himself that the Gödelian formula, although unprovable-in-the-system, was nonetheless – in fact, for that very reason – true." [Luc63]

point out the fundamental constraints of doing it in through mechanist or algorithmic methodology. This leads some to believe quantum mechanics may play an important role [Pen90, Ham03]. By formalizing what could not be done, Gödel had also demarcated what *would* be computable algorithmically. His work on self-referential formulae are a basis of a class of general problem solvers, later coined "Gödel machines" by Jürgen Schmidhuber. The generality of this problem-solving is an important aspect of general fluid intelligence and Gödel machines fare well in this respect. But since they are not biologically inspired we will not treat them more extensively.

Gödel machines or universal proof searchers can be digitally executed much faster than a human mind could perform the same operations. Schmidhuber claims [Sch06] that his Gödel machines execute O()-optimal, and can solve any solvable problem. But they are not exempt of Gödel's fundamental limitation [Sch06, p. 5].

Gödel's proof is not really debated, but the controversy about the many interpretations of Gödel's theorems shows that no view is conclusive agreed upon. Perhaps practical evidence can be gathered by evolving intelligence via biological metaphors.

4.1.4 John Searle's Chinese room experiment

John Searle disputed the notion that the brain is a computer at the fundamental level. He proposed the Chinese Room thought experiment. In short, it is a closed room with a person in it that knows no Chinese. He receives a message which he considers non-informational 'data', but he applies instructions that are provided on cards in the room. The instructions allow him to generate an answer that appears to come from someone that understands Chinese. Still, the man is not consciously aware of what is being discussed. While the man, the room nor the cards can be said to be 'aware' of it, there is no awareness in play. Similarly, an algorithmic machine would never be aware, or 'understand' what it processes. [Sea80]

4.2 Open ended evolution

> "[E]volution should be free to explore the possibilities without the burden of human "guidance"."
>
> `Alastair Channon` [Cha96]

4.2.1 Limitations of supervised learning

Evolutionary algorithms allow us to evolve a system to achieve a selected goal, without explicitly stating *how* to reach it. Specifying *what* to evolve towards, indirectly by specifying the fitness function, is an example of supervised learning. This capability is very appealing, but not without problems

[KLM96]. The fundamental problem of supervised learning used in an attempt to evolve complex behaviors, is that it is limited by the insight and creativity of the supervisor. A practical problem is that it intensively requires human involvement.

The situation is similar to programming a system to react intelligently. But then it is actually the programmer's intelligence that the system exhibits, and not its own. The extent of intelligent behavior is limited by the intelligence of the programmer. Many case based reasoning systems and knowledge-bases have been created, however most can only answer very narrow and specific questions. Common-sense, which appears in natural intelligence, is hard to recreate [Min06, Ch. 6]. Perhaps slowly, progress is made in reasoning with incomplete, inconsistent, imprecise and/or uncertain qualitative data [DP96, LS04]. As we have learned in section 3.3.2 (page 27), intelligence has a general aspect to it. And when it is entirely pre-programmed, it is more aptly called an instinctive than intelligent.

4.2.2 Evolutionary emergence

So far, we have been unable to exactly specify what intelligence is. In order to achieve emergence[3] of intelligence, Alastair Channon argues for the withdrawal of the traditional 'fitness function' based Genetic Algorithm (see figure 4.1a). We are unable to determine such a function that produces intelligence. In 'The Evolutionary Emergence route to Artificial Intelligence' [Cha96] Channon outlines evolution of a virtual world in which co-existence, interactions of species, acts as the selective force (see figure 4.1b). In such a world we can leave the direction of evolution open.

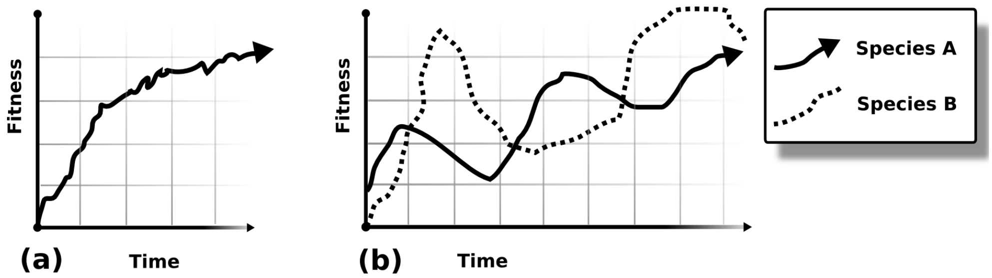

Figure 4.1: **(a)** Optimization **(b)** Co-evolution. Adapted from [Cha96].

His view is shared with some others, including Thomas Ray. Succinctly put, this is what happens in the Tierran world: "[O]nce the memory is filled with creatures, the creatures themselves become a prominent feature of the environment. Now evolution also discovers ways for creatures to exploit one another, and to defend against such exploitation." [Ray94, p. 245]

[3]Very simply put, with emergence more comes out of it than is programmed into it.

Channon concludes that the 'Selection' principle of Darwin's theory is often misinterpreted and overemphasized. Reminding us that the theory is one of local change and adaptation, not of optimization along an absolute scale of fitness. 'Selection' is a mere abstraction of probability of inheritance of any property.

Chapter 5

Conclusions

Nature has evolved an abundance of valuable and inspirational examples for us. It has even given birth to intelligent species. The hard-to-grasp concept of intelligence is valuable to us and deserves further exploration. Abilities such as abstraction, learning and dealing with novelty are important to intelligence. The brain processes underlying intelligence are still little understood [Got97, p. 14]. Luckily, complex and often complicated phenomena are found to have underlying principles that are not very complicated by themselves. In this thesis, **(1)** the principles for this evolutionary process were clarified and **(2)** evaluated for artificially application. **(3)** Finally we address whether intelligence can be said to emerge from the artificial systems.

Ad 1) General principles thought to be essential are evolution, evolvable structures (substrates) and interaction with a rich and challenging environment. Specifically, neuronal structures in the form of a brain have been essential to natural evolution of intelligence. The subsystems of the brain are modularized, but still highly coupled (semi-autonomous nuclei).

Ad 2) Both neuronal structures and evolution have been implemented artificially and have been combined, referred to as Artificial Neural Networks (ANNs) and Evolutionary Algorithms (EAs). As in biology, evolutionary and neuronal-inspired models are a strong combination.

Two experimental implementations are discussed and related to the theory. Experimenting with the evolution of shape (morphology) and control structure of virtual creatures results in some interesting observations and illustrates the problems one might encounter when designing selection functions. The Babybot experiment is quite unique as it takes a development approach to robotics, even while development is crucial for humans and since human like behavior is the premise of AI. In contrast with Goldstein and Godoy's creatures, the Babybot adds an interaction which is important

to more complex species, and essential to emergence of qualities such as intelligence: development.

Implementation challenges and issues are discussed, scaling and interconnection problems. Possible solutions are use of FPGA, aVLSI, neuromorphic engineering, optic-holographic and molecular computing devices.

General problem solving abilities are important to achieving artificial intelligence.

Ad 3) Do the implementations have what is needed for intelligence to emerge? Will intelligence eventually arise? A theoretical computer science perspective and the view of Alastair Channon are presented.

Gödel's incompleteness theorem and Searle's 'Chinese room' experiment are introduced. Channon argues that since we're unable to specify precisely what intelligence is, we should not expect it to emerge when using the fitness function in the traditional sense. Instead, Channon and others propose a co-evolution based approach.

Due to the complex nature of the topic, I must apologize to have pursued an approach often called eclectic[1]. Hopefully I have done this well.

[1]Selecting; choosing (what is true or excellent in doctrines, opinions, etc.) from various sources or systems; as, an eclectic philosopher.
Source: Webster's Revised Unabridged Dictionary (1913)

Appendix A

The Turing test

For the assessment of intelligence Alan Turing has devised a formal test which is now widely known as the Turing Test [Tur50]. This test was proposed as a supplement to the philosophically charged question "can machines think?", because that appeared to be very hard to answer.

A.1 Design testing procedure

It is based on an imitation game in which a man (A) a woman (B) and an examiner[1] (C) communicate via (typewritten) text messages. The objective of the examiner is to decide which is which by constructing smart questions. A computer is supposed to take the place of the man (A). The objective of (A) and (B) is to convince the examiner that it/she is the human person. According to the proposal the computer can be ascribed intelligence when it succeeds in convincing the examiner that it is the human.

Because of some objections as to the relevance of this question, you often find the additional requirements that: First, (C) should be given adequate time. Also (B) should be well capable of expressing herself, since someone can be intelligent without being able to express this well enough over a 'chat session'.

A.2 Relevance and objections

Alan Turing believed that some day a machine would be able to pass the Turing test (repeatedly to rule out luck or flaws of the interrogator or witness). It must be noted that AI is *useful* regardless of failing on this test. Paradoxically, it is currently also beneficial that some 'hard AI problem' are currently unsolved. It allows for alleviating automated abuse[2].

[1] Originally called the interrogator.

[2] Luis von Ahn has formalized the CAPTCHA, "Completely Automated Public Turing test to tell Computers and Humans Apart", which is in use for preventing unwanted

In advance, along with the proposal, Turing has addressed several objections that people may have:

- Theological objection: "God has endowed only humans with the gift of a soul and to be able to think.". Turing replies that God could create such a machine if He wishes so.
- 'Heads in the Sand' Objection: 'The consequences of machines thinking would be too dreadful. Let us hope and believe that they cannot do so.'.
- Mathematical objection: See below and the computability of thought (section 4.1). Still a topic of debate.
- Lady Lovelace's Objection: "[this machine] can do whatever we know how to order it to perform". Turing argues that we may one day know how to make it perform well enough to pass his test.
- Argument from Continuity in the Nervous System. See [Tur50]
- Informality of Behavior argument. [Ibid.]
- The Argument from Extra-Sensory Perception. [Ibid.]
- Consciousness. Still a topic of debate. Also by Penrose [Pen90, Pen96] and [Ibid.]
- Arguments from Various Disabilities. [Ibid.]

The objection of Penrose [Pen90] is mostly based on Searle's 'Chinese room' argument and Gödel's mathematical argument. Turing defends his opinion, but inconclusively: "Those who hold to the mathematical argument would, I think, mostly be willing to accept the imitation game as a basis for discussion.".

The test assesses AI in the category of "human intelligence" (see table 1.1 in section 1.3.2). It is not so clear with the 'thinks like' / 'acts like' classification. When the machine acts like a human and if it passes the test there would be no way to tell to what degree it thinks like a human.

So far, the test has not been passed when all requirements were applied. This illustrates the difficulty of solving 'hard AI problems'. It appears that Alan Turing realized this, he concluded his proposal as follows.

"We can only see a short distance ahead, but we can see plenty there that needs to be done."

Perhaps one day, when plenty of work is finished, the machine would succeed. And when it does, what will the next challenge be? (see figure A.1)

automatic usage of resources (e.g. subscription to e-mail services to be able to send SPAM).

Figure A.1: Turing Test 2.0. Courtesy of xkcd.com (CC ShareAlike license)

Bibliography

[Abe98] Abelard. the turing test and intelligence. Unpublished. From: http://www.abelard.org/turing/tur-hi.htm, 1998.

[Aza00] Farooq Azam. *Biologically Inspired Modular Neural Networks*. PhD thesis, Electrical and Computer Engineering, Virgina Tech, Virginia, USA, May 2000. http://scholar.lib.vt.edu/theses/available/etd-06092000-12150028/unrestricted/etd.pdf.

[BK92] E.J.W. Boers and H. Kuiper. Biological metaphors and the design of modular artificial neural network. Master's thesis, Departments of Computer Science and Experimental Psychology. Universiteit van Leiden, The Netherlands, 1992.

[Cal98] William H. Calvin. Competing for consciousness: A darwinian mechanism at an appropriate level of explanation. *Journal of Consciousness Studies*, 5(4):389–404, November 1998.

[Cal04] William H. Calvin. *A Brief History of The Mind*. Oxford University Press, UK, 2004. http://WilliamCalvin.com/BHM.

[CdZR93] A. Chiabrera, E. di Zitti, and D. Ricci. Biological paradigms of molecular electronics. 1993.

[CGC+98] William H. Calvin, Stephen Jay Gould, Yves Coppens, Ian Tattersall, and Luca Cavalli-Sforza. The emergence of intelligence. *Scientific American Presents*, 9(4):44–51, November 1998.

[Cha96] A. Channon. The evolutionary emergence route to artificial intelligence. Master's thesis, University of Sussex, October 1996.

[Cha01] A. Channon. *Evolutionary Emergence: The Struggle for Existence in Artificial Biota*. PhD thesis, University of Southampton, November 2001.

[Cla01] Andy Clark. Reasons, robots and the extended mind. *Mind and Language*, 16(2):121–145, 2001.

[dGK02] Hugo de Garis and Michael Korkin. The CAM-brain machine (CBM) – an FPGA based hardware tool which evolves a 1000 neuron net circuit module in seconds and updates a 75 million neuron artificial brain for real time robot control, 2002.

[dGKG+98] Hugo de Garis, Michael Korkin, Felix Gers, Norberto Eiji Nawa, and Michael Hough. "CAM-brain" ATR's artificial brain project - an overview, 1998.

[DP96] D. Dubois and H. Prade. Approximate and commonsense reasoning: From theory to practice. *Lecutre Notes in Computer Science*, 1079:19–??, 1996.

[EA03] Chris Eliasmith and Charles H. Anderson. *Neural engineering: computation, representation, and dynamics in neurobiological systems.* The MIT Press, Cambridge, Massachusetts, 2003.

[Elm05] J.L. Elman. Connectionists models of cognitive development: where next? *TRENDS in Cognitive Sciences*, 9(3), March 2005.

[ERR94] A.E. Eiben, P-E Raué, and Zs. Ruttkay. Genetic algorithms with multi-parent recombination. In *Proceedings of the 3rd Conference on Parallel Problem Solving from Nature*, LNCS Series. Springer-Verlag, 1994.

[Ese06] Deniz Eseryel. *Expert Conceptualizations of the domain of instructional design: An investigative study on the deep assesment methodology for complex problem-solving outcomes.* PhD thesis, Graduate School of Syracuse University, 2006.

[Fal05] Dean Falk. Constraints on brain size: The radiator hypothesis. In Todd M. Preuss and Jon H. Kaas, editors, *The Evolution of Primate Nervous Systems*, volume 4 of *Evolution of Nervous Systems*. Elsevier - Academic Press., 2005.

[Fau94] Laurence Fausett. *Fundamentals of Neural Networks: architectures, algorithms and applications.* Prentice Hall; US Ed edition, January 1994.

[For93] Stephanie Forrest, editor. *Proceedings of the Fifth International Conference on Genetic Algorithms*, San Mateo, CA, July 17-21 1993. University of Illinois at Urbana Champaign, Morgan Kaufmann.

[Gea05] David C. Geary. *Evolutionary Perspectives on Human Development*, chapter Evolution and Cognitive Development, pages 99–134. Sage Publications, 2^{nd} edition, 2005.

[GG07] Felipe Portavales Goldstein and Andy Godoy. Evolução de morfologia e controle de criaturas virtuais / evolution of morphology and control of virtual creatures. Retrieved August 2007.

[Got97] Linda S. Gottfredson. Mainstream science on intelligence: An editorial with 52 signatories, history, and bibliography. *Intelligence*, 24(1):13–23, 1997.

[Gre98] R. L. Gregory. The oxford companion to the mind. 1998.

[GW93] Scott V. Gordon and Darrell Whitley. Serial and parallel genetic algorithms as function optimizers. In [For93], pages 177–183, 1993.

[Ham03] Stuart R. Hameroff. *Ultimate Computating: Biomolecular Consciousness and NanoTechnology.* Elsevier Science Publishers B.V., Noord-Holland, Original 1987, reprint 2003.

[Haz02] Michiel Hazewinkel, editor. *Encyclopaedia of Mathematics.* Springer-Verlag Berlin Heidelberg New York, 2002.

[Hey96] Francis Heylighen. *The Evolution of Complexity*, chapter The Growth of Structural and Functional Complexity during Evolution. Kluwer Academic Publishers, 1996.

[Hol75] John H. Holland. *Adaptation in Natural and Artificial Systems.* The University of Michigan Press, Ann Arbor, 1975.

[Hop82] J. Hopfield. Neural networks and physical systems with emergent collective properties. In *Proceedings of the National Academy of Sciences USA*, volume 79, pages 2554–2558, 1982.

[Hut04] Marcus Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability.* Springer, Berlin, 2004. 300 pages, http://www.idsia.ch/~marcus/ai/uaibook.htm.

[JD00] et al. John Duncan. A neural basis for general intelligence. *Science*, 289(457), 2000.

[Jen99] Arthur R. Jensen. The g factor: the science of mental ability. *Psycoloquy*, 10(23), 1999.

[KLM96] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. 4:247–285, 1996.

[KN93] A. J. Kurdila and F. J. Narcowich. Sufficient conditions for penalty formulation methods in analytical dynamics. *Computational Mechanics*, 12:81–96, 1993.

[Kur90] Ray Kurzweil. *The Age of Intelligent Machines*. 1990.

[Kur99] Raymond Kurzweil. *The Age of Spiritual Machines*, chapter 6. Building New Brains. Kluwer Academic Publishers, 1999.

[KW01] Bryan Kolb and Ian Whishaw. *An introduction to Brain and Behaviour*. Worth Publishers, New York, 2001.

[LH06] S. Legg and M. Hutter. A collection of definitions of intelligence, 2006. From: http://www.vetta.org/shane/intelligence.html.

[LH07] S. Legg and M. Hutter. A collection of definitions of intelligence. In B. Goertzel and P. Wang, editors, *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms*, volume 157 of *Frontiers in Artificial Intelligence and Applications*, pages 17–24, Amsterdam, NL, 2007. IOS Press.

[Lip97] R. P. Lippmann. Speech recognition by machines and humans. *Speech Communication*, 22:1–15, 1997.

[LS04] H Liu and P Singh. Conceptnet – a practical commonsense reasoning tool-kit. *BT Technology Journal*, 22(4):211–226, 2004.

[Luc63] J. R. Lucas. Minds, machines and gödel. 1963. First published in Philosophy, XXXVI, 1961, pp.(112)-(127); reprinted in The Modeling of Mind, Kenneth M.Sayre and Frederick J.Crosson, eds., Notre Dame Press, 1963, pp.[269]-[270]; and Minds and Machines, ed. Alan Ross Anderson, Prentice-Hall, 1954, pp.43-59.

[McD05] Michael A. McDaniel. Big-brained people are smarter: A meta-analysis of the relationship between in vivo brain volume and intelligence. *Intelligence*, 33:337–346, 2005.

[Min06] Marvin Minsky. *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*, chapter 6. Simon and Schuster, November 2006.

[Moo65] Gordon E. Moore. Crammnig more components onto intergrated circuits. 38(8), 1965.

[MPMS00a] Giorgio Metta, Francesco Panerai, Riccardo Manzotti, and Giulio Sandini. Babybot: an artificial developing robotic agent. 2000.

[MPMS00b] Giorgio Metta, Francesco Panerai, Riccardo Manzotti, and Giulio Sandini. Development: Is it the right way towards humanoid robotics? *IAS-6*, July 2000.

[MSS95] John Maynard Smith and Eörs Száthmary. *The Major Transitions in Evolution.* Freeman, Oxford, 1995.

[MT95] Geoffrey F. Miller and Peter M. Todd. The role of mate choice in biocomputation: Sexual selection as a process of search, optimization, and diversification. In Wolfgang Banzhaf and Frank H. Eeckman, editors, *Evolution and Biocomputation: Computational Models of Evolution*, pages 169–204. Springer, Berlin, 1995.

[Nol01] David D. Nolte. *Mind at Light Speed: a new kind of intelligence.* The Free Press, Simon & Schuster, Inc., 2001.

[Opa38] A.I. Oparin. *A. I. (1938). The Origin of Life. Translated by Sergius Margulis.* Macmillan, New York, 1938.

[Pen90] Roger Penrose. *De Nieuwe Geest van de Keizer: Over computers, de menselijke geest en de wetten van de natuurkunde (Dutch translation of "The Emperor's New Mind").* Prometheus, Amsterdam, 1990.

[Pen96] Roger Penrose. Beyond the doubting of a shadow: A reply to commentaries on shadows of the mind. *PSYCHE*, 2(23), 1996.

[Pia36] Jean Piaget. *The origin of intelligence in children.* Interntational University Press, Inc. and Routledge & Kegan Paul Ltd., London, 1936. M. Cook, Trans.

[PWB87] D. Psaltis, K. Wagner, and D. Brady. Learning in optical neural computers. *IEEE*, Vol. III:pp. 549–555., 1987.

[QS82] S. Quartz and T. Sejnowski. The neural basis of cognitive development: A constructivist manifesto, 1982.

[Ram98] Ramachandran. Optoelectronic implementation of neural networks. *Resonance*, pages 45–55, September 1998.

[Ray94] Thomas S. Ray. Evolution, complexity, entropy, and artificial reality. *Physica*, (D 75):239–263, 1994.

[Ray99] Thomas S. Ray. Tierra tutorial. Online: www.his.atr.jp/~ray/pubs/tutorial/Tutorial.ppt, July 1999.

[Reb95] Arthur S. Reber. *Dictionary of Psychology.* Penguin Books, United Kingdom, 2^{nd} edition, 1995.

[Ril87] H. Norton Riley. The von neumann architecture of computer systems. California State Polytechnic University, September 1987.

[RMK06] Tarassenko R. Morris and M Kenward, editors. *Cognitive Systems: Information Processing meets brain science.* Elsevier, Amsterdam, The Netherlands, 2006.

[RN95] S. Russell and P. Norvig. Artificial intelligence: A modern approach. Prentice Hall Series in Artificial Intelligence, Englewood Cliffs, New Jersey, 1995.

[Roj96] Raúl Rojas. *Neural Networks - A Systematic Introduction.* Springer-Verlag, Berlin, New-York, 1996.

[Sch06] Jürgen Schmidhuber. Gödel machines: Self-referential universal problem solvers making provably optimal self-improvements. *TR IDISIA*, 19(3), December 2006.

[SD02] Brian Lau Garrett B. Stanley and Yang Dan. Computational subunits of visual cortical neurons revealed by artificial neural networks. *PNAS*, 99(13):8974–8979, June 2002.

[Sea80] John R. Searle. Minds, brains, and programs. *Behavioral and Brain Sciences*, 3 (3):417–457, 1980.

[Sha02] David R. Shaffer. *Developmental Psychology: Childhood and Adolescence.* Wadsworth, Thomson Learning, 6th edition, 2002.

[Swe03] Latanya Sweeney. That's AI?: A history and critique of the field. July 2003.

[SZ98] Noel E. Sharkey and Tom Ziemkey. A consideration of the biological and psychological foundations of autonomous robotics. *Connection Science*, 10:(3-4):361–391, December 1998.

[Tur50] Alan M. Turing. Computing machinery and intelligence. *Mind*, LIX(236):433–460, Oct 1950.

[Vro89] P.A. Vroon. *Tranen van de krokodil: over de te snelle evolutie van onze hersenen.* Ambo, Baarn, 1989.

[Wan06] Pei Wang. *Rigid Flexibility: The Logic of Intelligence.* Springer, Dordrecht, The Netherlands, 2006.

[Wik07a] WikiPedia. Evolution (disambiguation), http://en.wikipedia.org/wiki/evolution_(disambiguation) - retrieved on october 2007. Online, October 2007.

[Wik07b] WikiPedia. Sigmoid function, http://en.wikipedia.org/wiki/sigmoid_function - retrieved on august 2007. Online, August 2007.

[WVL00] J. C. Wickett, P. A. Vernon, and D. H. Lee. Relationships between factors of intelligence and brain volume. *Personality and Individual Differences*, 29:1095–1122, 2000.

[ZS03] J. Zhu and P. Sutton. FPGA implementations of neural networks - a survey of a decade of progress. Unpublished, 2003.